

26/prts.

1 422 Rec'd PCT/PTO 22 SEP 2000.

OPTICAL COMMUNICATIONS NETWORK

The present invention relates to an optical communications network, and in particular to the regeneration of optical packets carried on such a network.

- It is well known that optical fibre has a huge potential information-carrying capacity. For example, by utilising the entire gain bandwidth of erbium-doped optical amplifiers, a single fibre could carry more than 2 Tbit/s. However in the majority of telecommunications systems in commercial use currently, the information is carried over fibre in the form of an optical signal at a single wavelength. The data transmission bandwidth of the fibre is therefore limited by the electrical bandwidth of the transmitter and receiver, and this means that only a tiny fraction (a maximum of about 1%) of the potential bandwidth-carrying capacity of the fibre is being usefully exploited. There is therefore much interest currently in developing methods for increasing the transmission rate for point-to-point fibre links. One method is wavelength-division multiplexing (WDM), in which several data channels, at different wavelengths, are carried simultaneously on the same fibre. An alternative method for increasing the rate of information that can be carried on fibre is to use optical time-division multiplexing (OTDM) in which several data channels are multiplexed in the form of bit-interleaved return-to-zero (RZ) optical pulse trains.
- The WDM approach to photonic networking has some very attractive advantages: in addition to the relative simplicity and commercial availability of the devices needed, WDM networks can be created in a wide variety of architectures with great flexibility (the main restriction being merely that any pair of photonic transmission paths cannot use the same wavelength on a shared fibre link). An advantage of WDM networks is that they can, in principle, support 'signal transparency', i.e. data signals can be carried using any modulation format. However, this implies that, in effect, WDM photonic networks are based on 'analogue' transmission. As a result it is not possible for digital signal regeneration techniques in the optical domain, to be used. The inability to perform signal regeneration in the optical domain leads to practical scaling limitations for WDM networks due to noise accumulation from optical amplifiers, crosstalk and nonlinearity. These factors restrict the number of network switching nodes through which signals can pass without fatal degradation. Currently, in reported laboratory experiments the maximum number of WDM switching nodes through which a

signal can pass without regeneration is limited to around 10, which is a significant restriction in architecture and scalability. A feasible, though costly, solution currently being advocated by some equipment vendors is to sacrifice transparency, standardise the transmission format, and regenerate each wavelength channel
5 individually at the outputs of WDM cross-connects. In effect, this is a hybrid arrangement using analogue switching together with channel-by-channel digital regeneration.

In the OTDM approach to photonic networking, the signals are carried in 'digital' format in the form of RZ optical pulses, allowing the use of digital signal
10 regeneration techniques in the optical domain such as 3R (Re-amplify, Re-time and Re-shape) regeneration [Lucek J K and Smith K, Optics Letters, 18, 1226-28 (1993)] or soliton-control techniques [Ellis A D, Widdowson T, Electronics Letters, 31, 1171-72 (1995)]. These techniques can maintain the integrity of the signals as they pass through a very large number of nodes. For example, Ellis and Widdowson [Ellis A
15 D, Widdowson T, Electronics Letters, 31, 1171-72 (1995)] have made a laboratory demonstration of error-free transmission of signals through an OTDM network consisting of 690 nodes in concatenation. Despite this impressive potential for scalability, however, the OTDM approach to photonic networking suffers from severe restrictions in the network architecture that can be used. This results from
20 the need to maintain proper bit-level synchronism between all the signal sources, demultiplexers and channel add/drop multiplexers throughout the network. In complex architectures, such as one involving the merging of signal streams emanating from several widely-separated sources, fluctuations in the arrival time of signals (due to environmental effects acting on the fibres such as temperature
25 change and mechanical strain) cannot be adequately controlled or compensated in a continuous uninterrupted fashion. This is because of the restricted range of variable optical delay lines, the limited frequency response of control systems due to the physical time of flight of signals over extended distances, and also insufficient degrees of freedom.

30 According to a first aspect of the present invention, there is provided a method of operating a node in an optical communications network including
a) receiving at the node an optical packet; and

b) generating from the said optical packet received at the said node a regenerated optical packet having a phase determined by a local clock source and independent of the phase of the said packet received at the node.

Throughout this document, the term 'packet' is used to mean a fixed-length or variable-length string of bits which may be routed through a network in a variety of different ways, including self-routing, store-and-forward packet routing, scheduled switched routing and circuit switching

The present invention provides a new approach to operating an optical communications network which for the first time makes it possible for an optical packet network to be scaled almost without limit. This is achieved by carrying out regeneration of optical packets in a manner which is asynchronous at the bit-level relative to the packet source. This then allows digital signal regeneration to be carried out, facilitating the transmission of packets over large distances, without the extent and architecture of the network being constrained by the need to transmit a global bit-level timing signal.

Preferably the step of generating a regenerated optical packet includes gating, using the received optical packet, an optical clock signal from the local clock source. Preferably this includes passing the optical clock signal through each of a plurality of gate means;

applying the received optical packet as a control signal to each of the plurality of gate means with different delays of a fraction of a bit period relative to the optical clock signal input to the gate means; and

selecting the output of one of the plurality of gate means to provide the regenerated optical packet. Preferably the difference in delays is equal to T/k where T is the bit period and k is the number of optical gate means. Preferably the width W of the gate window is not less than T/k and not more than T , where T is the bit period and k is the number of optical gates.

Alternatively, as further described below, a single gate means may be used in conjunction with means to shift the phase of the incoming packet to match that of the local free-running optical clock.

It is particularly advantageous to use a free-running local optical clock which is gated by the incoming packet. The use of a free-running optical clock, as opposed to one which is phase-locked to an incoming signal, allows the source to be a high quality device such as a passively mode-locked laser. A further

advantage is that the clock may be used to control further optical processing devices at the local node in order, for example, to process header data carried with the packet.

It is found to be particularly effective to use a plurality of optical gates which are gated by the optical packet with different respective delays. The delay in question is that of the optical packet *relative* to the signal from the optical clock source. In practice the different relative delays may be achieved by applying different delays to the optical clock signal, or applying different delays to the optical packet. By using a number of gates in this way, and selecting one of their outputs, it is possible to recover an appropriately regenerated optical signal whatever the phase of the incoming optical packet. The optical gate means may comprise a number of distinct physical devices, such as those described in further detail below. Alternatively the plurality of optical gate means might comprise a single device arranged to gate a plurality of distinct optical signals distinguished, e.g. by their polarisation or wavelength, and references in the description and claims to a number of gate means are to be construed accordingly.

Preferably the method includes making a measurement of a parameter of an optical signal output from the gate means, and selecting the output of one of the plurality of gates to provide the regenerated optical packet depending on the results of the said measurement. The parameter may be a measure of the total energy of the output signal, and the selection may be made by comparing the energies of the signals from different optical gate means. Other parameters may also be used to make the selection. For example, the bit error level may be measured, and the output selected which has the minimum bit error level.

According to a second aspect of the present invention, there is provided a method of operating a communications network comprising a plurality nodes interconnected by an optical transmission medium, the method including:

transmitting an optical packet and at a network node, receiving the said packet and generating from the said packet a regenerated optical packet having a phase determined by a local optical clock source and independent of the phase of the said packet received at the network node.

According to a third aspect of the present invention, there is provided a regenerator for optical packets comprising:

an input for receiving an optical packet;

a local optical clock source comprising a free-running local oscillator;
means for generating from the optical packet and from a clock signal from
the local optical clock source a regenerated optical packet independent in phase
from the optical packet received at the input.

- 5 Systems embodying the present invention will now be described in further detail, by way of example only, and will be contrasted with the prior art, with reference to the accompanying drawings, in which:
- Figure 1: shows the basic concept for 3R regeneration of a digital data stream consisting of a RZ pulse train encoded by on-off modulation
- 10 Figure 2: shows the method of 3R regeneration used for OTDM systems in which the bit stream is continuous
- Figure 3: shows a method of partial regeneration of optical packets
- Figure 4: a simplified outline diagram showing a generator of optical packets and a bit-asynchronous packet regenerator
- 15 Figure 5: a dual-gate bit-asynchronous packet regenerator
- Figure 6: a sequence of timing diagrams that illustrate the operation of the dual-gate regenerator
- Figure 7: diagram showing the probability density function of the arrival time of the i th data bit at the gates, relative to the clock signal input to gate G1
- 20 Figure 8: a plot of the expression (12) against the phase angle θ , taking the values $\sigma = 0.018T$ and $W = 0.75T$
- Figure 9: shows a plot of the bit-error probability B against the phase angle θ , calculated according to (12), using an ideal mechanism to select the output gate i , and taking $W/T = 0.75$ and $\sigma/T = 0.018$ and 0.03
- 25 Figure 10 shows a plot of the maximum bit-error probability B according to (12), plotted versus the bit-arrival jitter σ , and taking $W/T = 0.75$.
- Figure 11: is a plot of the expected bit-error probability for a given packet, per regenerator passed, according to (15), plotted against the rms jitter σ , with $W/T = 0.65, 0.75, 0.85$.
- 30 Figure 12: shows a plot of the bit-error probability B against the phase angle θ , calculated according to (18), taking $W/T = 0.75$, $\sigma/T = 0.018$ and $\sigma_c/T = 0.001$
- Figure 13 shows a plot of the maximum value of the bit-error probability B for the phase angle θ anywhere in the range $0 \leq \theta < 2\pi$, calculated according to (18), plotted versus σ_c/T , with $W/T = 0.75$ and $\sigma/T = 0.001, 0.018$ and 0.03 .

Figure 14: Diagram of a bit-asynchronous quad-gate packet regenerator

Figure 15: plot of bit-error probability according to (19) plotted against the phase angle θ , taking $\sigma = 0.036T$ and $W = 0.75T$

Figure 16: a plot of the bit-error probability B against the phase angle θ ,
 5 calculated according to (12), using an ideal mechanism to select the output gate /
 (such as the strategy (20) executed by ideal circuits without systematic or random
 errors) and taking $\sigma = 0.036T$ and $W = 0.75T$

Figure 17: Plot of the bit-error probability B against the phase angle θ , calculated
 according to (21), and assuming that $W/T = 0.75$, $\sigma/T = 0.036$ and
 10 $\sigma_c/T = 0.001$ and 0.052 ."

Figure 18: Plot of the maximum value of the bit-error probability B for any phase
 angle θ in the range $0 \leq \theta < 2\pi$, calculated according to (21), plotted versus
 W/T , with $\sigma/T = 0.036$, and $\sigma_c/T = 0.001$ and 0.05 ."

Figure 19: show values for the probability that a packet suffers 'slippage' (i.e. a
 15 time displacement greater than a specified acceptable limit $L T$), according to (22),
 plotted versus N , the number of regenerators passed, for various values of $L T$.

Figure 20: shows a version of the two-gate regenerator which uses only one
 optical gating device to gate simultaneously two independent clock signals which
 are distinguishable by their different states of polarisation.

20 Figure 21: Outline diagram of a bit-asynchronous packet regenerator using a single
 gate to modulate the output of the local source.

Figure 22: Example embodiment of the bit-asynchronous packet regenerator using
 a single gate to modulate the output of the local source.

Figure 23: Example layout of a portion of a network comprising switching nodes
 25 containing routing switches (RS) and bit-asynchronous regenerators (AR), and links
 between the switching nodes containing bit-synchronous regenerators (SR).

Figure 24: Outline diagram of an alternative arrangement of the bit-asynchronous
 packet regenerator using a single gate to modulate the output of the local source.

30 Figure 25: A further example embodiment of the bit-asynchronous packet
 regenerator using a single gate to modulate the output of the local source

Figure 26 shows a schematic of an optical packet network.

In an ultrafast optical packet network information is transported across a network
 in the form of fixed-length bursts (i.e. cells or fixed-length packets) of RZ optical

DOCUMENT NUMBER

- pulses that are encoded with payload data and control information (such as the address of the packet destination). The network may constitute, for example, the core network of a national data/telephony network, or a local area network connecting a number of computing systems, or part of the fabric of a
- 5 communications switch, or may provide the connection between processors in a multi processor computer. Examples of suitable topologies for such a network are described in the present applicant's co-pending European patent application 97307224.2, the contents of which are incorporated herein by reference. The use of logical (header), rather than physical, addressing facilitates massive scalability.
- 10 Every transmission path in the network carries a continuous sequence of time slots, synchronised to a packet-level global clock, and each time slot accommodates at most one packet and an appropriate time guard band. This time guard band allows the transmission path to be switched for packet-by-packet routing and also allows continuous and endless synchronisation of the packet
- 15 streams and network nodes to the packet-level global clock. Thus the packet time slots are synchronised throughout the network. However, a crucial aspect of our approach is that synchronism between packets at the (picosecond) bit-level is not required. Therefore the portion of the time slot that accommodates the packet can be conveniently defined to be several bit periods longer than the packet duration.
- 20 This permits a certain amount of slop in the positioning of the packet within the time slot - in other words, the positioning of the packet within the time slot is not made with bit-level precision, and generally successive packets are not bit-synchronous.
- For such an asynchronous network to have the same near-infinite scalability as bit-
- 25 synchronous OTDM networks, it is necessary to use a bit-asynchronous digital optical packet regenerator, as is further described below.

Figure 1 shows a prior art approach to 3R regeneration of a digital data stream consisting of a RZ pulse train encoded by on-off modulation ('mark' represents a

30 bit value 1, 'space' represents 0). The incoming data bits from a distant source are used to modulate a continuous train of high-quality RZ pulses produced by a local source, thus regenerating the original data. The presence of a 'mark' in the incoming data stream causes the gate to open for a time of the order of the bit period, allowing a single pulse from the local source to pass through. In this way

the regenerated bits have the same pulse shape, spectral quality, amplitude and timing stability as the local source. The pulse repetition rate of this local source is the same as the bit rate of the incoming data. The key problem is designing such a regenerator is to ensure that the incoming data stream and the locally-generated
5 pulses are maintained in synchronism.

Figure 2 shows a prior art method used for OTDM systems in which the bit stream is continuous. A clock recovery circuit (which may be electronic or optical) derives a clock signal in synchronism with the incoming data bits, and this clock is used to synchronise the local pulse source.

- 10 The method shown in Figure 2 is less suitable for systems in which the incoming data is not continuous, but is in the form of bursts or packets which are mutually bit-asynchronous. One reason is that in this case the clock recovery must be performed on a packet-by-packet basis, and therefore the lock-up time of the
15 clock-recovery circuit must be a fraction of the packet duration in order not to waste bandwidth. However in ultrafast systems the need for a fast lock-up time tends to conflict with the requirement for a consistent high-quality pulse train throughout the duration of the packet. Another reason for the unsuitability of the approach shown in Figure 2 for ultrafast optical systems is that the best quality pulses are obtained from passively-mode-locked lasers and mode-locked ring lasers
20 which are not readily synchronisable to an external clock.

For some processes in optical packet networks, a self-synchronising approach has been proposed in which each incoming packet contains one or more synchronisation or 'marker' pulses which may be distinguishable by their wavelength, polarisation, amplitude, position, pulse spacing, etc. Figure 3 shows
25 the approach in which a synchronisation pulse is extracted from an incoming packet and then passively replicated to produce a regular burst or pulse pattern. This replicated pulse train can then be used in the processes of demultiplexing, packet header-address recognition. It may also be used in the process of partial regeneration as shown in Figure 3 . In this case the data bits in the incoming
30 packet are used to gate the train of replicated synchronisation pulses. However this is not full 3R regeneration: although amplitude and timing fluctuations are suppressed in the regenerated packet, the pulse shape and spectral quality are unchanged and therefore the pulse degradation suffered during transmission cannot be corrected.

In the usual approach to 3R regeneration, illustrated in Figure 2, the local pulse source is forced into frequency and phase synchronism with the incoming data stream. In the embodiments of the present invention described below, a different approach is used, in which the local pulse source is continuously free-running,

5 with the same nominal repetition frequency as the bit rate of the incoming data, but not necessarily in frequency and phase synchronism with the incoming packet. Instead the regenerated packet simply takes up the bit rate and phase of the free-running local source. This new approach allows the data packets to be bit-asynchronous, yet avoids the difficulty of recovering a high-quality clock signal on

10 a packet-by-packet basis. Also since the local source of RZ pulses is continuously free-running, it can be a high-quality source such as a passively-mode-locked laser or a mode-locked ring laser. A further benefit is that the regenerated packets are all in precise bit synchronism with the local source of RZ pulses, and therefore this source may be used as the continuous and regular source of pulses that is

15 essential for the various digital optical processing schemes, such as optical memory, registers, parity-counters, examples of which are described and claimed in the present applicant's copending applications number GB9726477.4, the contents of which are incorporated herein by reference. Such circuits find application in high-speed 'on-the-fly' processing of packets for routing, signalling,

20 error detection, etc.

3. Bit-asynchronous optical packet regeneration

3.1 Outline description

Figure 4 is a simplified outline diagram showing a generator of optical packets and a bit-asynchronous packet regenerator. The packet generator creates optical data

25 packets each with a fixed bit rate of $M_S f_S$. As shown in Figure 4, this packet generator could consist of a source of optical RZ pulses at a repetition frequency f_S Hz, whose output is modulated and multiplexed in a fashion similar to that used for OTDM (i.e. the output from the pulse source is split into M_S parallel paths which are individually encoded with data by on-off modulation at a rate f_S bit/s

30 and then recombined by bit-interleaving to form a packet of data bits with a composite rate of $M_S f_S$ bit/s). The source of pulses at repetition frequency f_S could, as shown in Figure 4, consist of an electronic microwave oscillator at f_S which drives an electrically-synchronised laser (such as a gain-switched laser or an

actively-mode-locked laser). Alternatively, it could be a continuously free-running optical pulse source, such as a passively-mode-locked laser or a mode-locked ring laser, whose nominal repetition frequency is set (for example, by tuning the laser cavity length). At the distant regenerator there is an independent continuously

- 5 free-running source of RZ optical pulses at a repetition frequency $M_R f_R$ Hz. This pulse source could, as shown in Figure 4, consist of an independent electronic microwave oscillator at a frequency f_R which drives an electrically-synchronised laser (such as a gain-switched laser or an actively-mode-locked laser), followed by a passive splitter and bit-interleaver to produce the required composite frequency
- 10 $M_R f_R$ Hz. Alternatively the source could be a continuously free-running optical source of pulses at the rate f_R followed by the $\times M_R$ splitter and interleaver, or indeed the free-running source could operate at the full repetition frequency $M_R f_R$ Hz directly without the need for $\times M_R$ multiplication. In any case, the bit rate $M_S f_S$ of the packet produced by the packet generator and the repetition frequency
- 15 $M_R f_R$ of the pulse source in the regenerator must lie close to the same nominal value $1/T$,

$$\text{i.e. } M_S f_S \approx M_R f_R \approx 1/T \quad (1)$$

where T is the nominal bit period in the optical packet. (For example, $T = 10$ ps, $1/T = 100$ Gbit/s, $M_S = 10$, $1/M_S T = 10$ Gbit/s.)

- 20 This novel approach in which the packet source and packet regenerator contain independent free-running pulse sources works because the optical data packets can be assumed to be of short duration, and therefore any phase slippage between the two pulse sources that occurs over the duration of a single packet is negligibly small. This can be illustrated by considering the question: how close must be the
- 25 bit rate of the packet and the repetition frequency of the pulse source at the regenerator?

At the inputs to the gate in the regenerator, the phase of the local clock pulses (input A in Figure 4) relative to the incoming packet data bits (control input C) as a function of time t is

$$30 \quad \theta(t) = 2\pi(M_R f_R - M_S f_S)t + \theta_0, \quad (2)$$

where θ_0 is a constant offset. The change in $\theta(t)$ over the duration of a packet of length n bits is

$$\Delta\theta = \theta(t + (n-1)T) - \theta(t) = 2\pi(M_R f_R - M_S f_S)(n-1)T. \quad (3)$$

This corresponds to a time slippage ΔT in the local clock relative to the packet data bits,

$$\Delta T = \Delta\theta \cdot T / 2\pi = (M_R f_R - M_S f_S)(n-1)T^2, \quad (4)$$

and this time slippage can be considered negligible if $|\Delta T| \ll T$. For example, with
 5 electrically-driven sources (as shown in Figure 4), $f = 10$ GHz, $M_S = M_R = 10$,
 $T = 10$ ps, $1/T = 100$ Gbit/s, $n = 1000$, and $f_R - f_S < 10$ kHz (< 1 in 10^6) results in,
 $\Delta T < 10^{-3} T = 10$ fs. This example shows that with realistic parameters, the time slippage over the duration of a packet can be negligibly small.

Throughout the remainder of this work it will be assumed that the bit rate of the
 10 incoming data packet and the repetition frequency of the pulse train at the A input to the gate can be considered to be exactly equal, and that the phase difference θ between them can be considered to be constant over the duration of the packet. Nevertheless, for any given packet, this phase difference is an arbitrary value in the range $0 \leq \theta < 2\pi$. Moreover, in a packet network, successive packets arriving
 15 at a regenerator may originate from different sources and accordingly their phases θ may be entirely uncorrelated. Therefore the bit-asynchronous packet regenerator must be designed to operate correctly on each packet independently, regardless of the phase difference θ . This is achieved in this invention by providing not just one gate (as shown in Figure 4) but two or more gates, each of
 20 which can correctly regenerate a packet for only a restricted range of θ , but which together span the complete range $0 \leq \theta < 2\pi$, and by providing means to select the output from one of the gates in each time slot so as to choose a correctly regenerated packet.

This is best explained and illustrated by examples. The next sub-section describes
 25 the structure and operation of a bit-asynchronous packet regenerator containing two gates, and the succeeding sub-sections describes regenerators with four gates.

3.2. Dual-gate bit-asynchronous optical packet regenerator

3.2.1 Principle of operation

30 Figure 5 shows a dual-gate bit-asynchronous packet regenerator. The data bits in the incoming packet are used to control the opening of two gates, G1 and G2. A data bit with value 1 ('mark') causes each of the two gates to open for a fixed time duration (the gate window), otherwise the gates remain closed. It is

preferable, though not essential, that the widths of the time window for gates G1 and G2 are equal. The output from the local clock (a continuous free-running source of optical RZ pulses at a repetition frequency nominally equal to the packet bit rate $1/T$) is applied to the inputs of the two gates, one of these inputs being delayed relative to the other by an amount $T/2$. Since the phase θ of the local clock pulses relative to the packet data bits has an arbitrary and unknown value in the range $0 \leq \theta < 2\pi$, it is necessary that the gate window widths are chosen so that, whatever the value of θ , the clock pulses will be correctly modulated by at least one of the two gates. In the case that the window widths for gates G1 and G2 are equal, the window width W must therefore lie in the range $T/2 < W < T$. The lower limit ensures that at least one clock pulse will be modulated at any value of θ , whilst the upper limit comes from the requirement that no more than one clock pulse may pass through the gate whilst the window is open. These upper and lower limits on W apply strictly in the case that the incoming data pulses and local clock pulses are sufficiently narrow that, on the time scale of a bit period, they may be represented by delta impulses. When finite pulse widths are taken into account the acceptable range of window widths is somewhat narrower than $T/2 < W < T$. Until sub-section 3.4.4 it will be assumed that the data pulses and local clock pulses are impulses, and for the remainder of sub-section 3.2 it is taken that $W = 3T/4$ for both gates.

We should note at this point, with reference to Figure 5, that there is an alternative and equally valid configuration in which the $T/2$ delay line is removed from the input port A of one of the gates, and placed instead at the control port C of one of the gates. The operation of the regenerator is very similar in this case, and the predicted performance described later is the same. Throughout the remainder of this description, the configuration will be assumed to be that in which the packet data bits are connected directly to the control ports of the gates, and the input ports have certain differential delays (as illustrated in Figure 5 for a dual-gate regenerator).

The optical gates may be implemented in different ways. For ultrafast operation, the gate could be a nonlinear optical device such as a fibre loop mirror (as described, for example, by Whitaker et al in Optics Letters, vol. 16, page 1840 (1991)), in which case the gate width is defined by selecting suitable fibre lengths, dispersion and birefringence. Alternatively a suitable ultrafast gating device based

on the nonlinearity in semiconductor optical amplifiers could be used (as described, for example, by Kang et al in the International Journal of High Speed Electronics and Systems, vol. 7, page 125 (1996)). In this case the gate width may be determined by the positioning of the amplifier in a Sagnac interferometer loop
5 arrangement, or the relative offset of two amplifiers in a Mach-Zehnder interferometer device. Another suitable ultrafast semiconductor-based device is the ultrafast nonlinear interferometer switch described by Hall and Rauschenbach (paper PD5, Proceedings of Conference on Optical Fiber Communication OFC'98, published by the Optical Society of America, February 1998), which has been
10 shown to operate at a speed of 100 Gbit/s. For operation at lower speeds an optoelectronic device such as an electroabsorption modulator could be used as the gate. In that case, the incoming packet data bits must first be received by a photodetector whose output is converted to a suitable short electrical pulse to drive the modulator, and the gate width is defined by the width and amplitude of
15 this electrical pulse. In this case, for correct operation it is necessary that the photodetector and associated electronics can fully resolve the data bits, which limits the packet data rate.

If a single optical gating device is used for the purpose of gating k distinct optical clock signals ($k > 1$) then, throughout the description and analysis in this work,
20 this will be regarded as k optical gates. The clock signals may be distinguishable, for example, by virtue of their states of polarisation, wavelength or intensity. For example, Figure 20 shows a version of the two-gate regenerator which uses only one optical gating device to gate simultaneously two independent clock signals which are distinguishable by their different states of polarisation.
25 Figure 6 shows a sequence of timing diagrams that illustrate the operation of the dual-gate regenerator. The packet data bits (an example sequence 11101 is shown) arrive at the control ports of the gates G1 and G2, and each 'mark' causes the gates to open for a time $3T/4$. The diagrams (i-iv) illustrate various values of θ , the phase of the local clock relative to the incoming packet data bits. It is
30 convenient to write the gate width W and phase angle θ as normalised quantities, thus: $\bar{W} = W/T$ and $\bar{\theta} = \theta/2\pi$. Diagram (i) illustrates the case $0 \leq \bar{\theta} \leq \bar{W} - 1/2$, in which the outputs from both gates G1 and G2 are clock pulses that have been correctly modulated by the incoming data bits. Diagram (ii) illustrates the case $\bar{W} - 1/2 < \bar{\theta} < 1/2$, in which the output from gate G1 (but not G2) are clock

pulses that have been correctly modulated by the incoming data bits. Diagram (iii) illustrates the case $1/2 \leq \bar{\theta} \leq \bar{W}$, in which again the outputs from both gates G1 and G2 are correctly modulated. Diagram (iv) illustrates the last possibility, $\bar{W} \leq \bar{\theta} < 1$, in which the output from gates G2 (but not G1) is correctly modulated.

- 5 Returning to Figure 5, the components shown to the right-hand side of the two gates are used to attempt to select in each time slot whichever gate output gives a regenerated packet with the minimum of bit errors. One technique, shown in Figure 5, is to make the selection in each time slot on the basis of a comparison of the total optical energy emerging from each gate, integrated over the duration of
- 10 the packet. If the phase angle θ is such that the output from a gate consists of correctly modulated clock pulses then the total optical energy measured at the output of the gate, integrated over the duration of the packet, will be maximum (in effect, it is a measure of the number of 'marks' appearing in the regenerated data packet). However if θ is such that the clock pulses arrive at the gate at a time
- 15 outside the gate window, then the energy transmitted by the gate will be zero or small. The circuit shown in Figure 5 therefore makes these energy measurements and the result of the comparison is used to set the optical switch S, (for example, an optoelectronic device such as a lithium niobate switch) which performs the physical selection. The detectors, D1 and D2, are followed by electronic
- 20 integrators I1 and I2, each of which provide a voltage proportional to the energy of the regenerated packet emerging from gates G1 and G2, respectively, in each time slot. The comparator C produces a digital output according to whether or not the signal from D1 exceeds that from D2. The global packet-level clock signal, synchronised to the time guard band between packets, is used to reset the
- 25 integrators and also to clock the D-type flip-flop DT. This ensures that the switch S changes over only during the guard band, so as to avoid corrupting a packet. Notice that the detectors, switch and associated electronics operate at the packet rate (not the data bit rate) with a response time on the order of the width of the time guard band (which may be on the time scale of ~1 ns).
- 30 An alternative method of selecting the most appropriate gate output in each time slot may be to perform a bit-error measurement on the whole or part of the packet that emerges from each gate. For example, a test pattern could be incorporated as part of each packet, and this pattern would be received and any bit errors

detected and counted in each time slot at the output of each gate. In a given time slot the output having zero or the least number of bit errors would be selected. The optical delays (labelled L and $L + \Delta T/2$ in Figure 5) between the outputs of the gates and the selection switch S are used to allow sufficient time for the 5 circuitry and switch S to operate before the packets arrive at the switch. Typically the delay L will be slightly less than one time slot in duration. Optionally, as shown in Figure 5, the optical delay between the output of gate G1 and the switch S may be made slightly longer (by an amount $T/2$) than the delay between the output of gate G2 and the switch. The purpose of doing this is to compensate for 10 the delay of $T/2$ at the input of gate G2, thus equalising the delay of both optical paths from the clock source to the output of the selector switch S. One benefit of doing this is that all the regenerated packets are then in precise bit synchronism with each other and with the local clock, and as mentioned earlier, the local clock may therefore be used as a continuous and regular source of pulses for use in 15 subsequent digital optical processing stages. A further important benefit of equalising the delay in this way is to reduce the problem of 'packet slippage' in a large network, as described in section 3.4.3.

There are two main causes of bit errors that may occur in the process of regenerating a packet using the bit-asynchronous regenerator. The first cause is 20 jitter in the arrival time of the incoming packet data bits. It is well known that in high-speed optical transmission systems, jitter in the arrival time of pulses arises from effects such as amplified spontaneous emission noise, the soliton self-frequency shift arising from the Raman effect, soliton short-range interactions, and the complex interplay of these various processes. If the time of arrival of a data 25 'mark' at the regenerator fluctuates relative to the local clock pulse train, the time position of the gate window opened by that data bit is shifted accordingly. This increases the probability that a clock pulse may fall outside the gate window and so may fail to be transmitted correctly as a 'mark' in the regenerated packet. The second main cause of bit errors is errors in the process used to select the output 30 from one of the gates in each time slot. For example, when the selection is made on the basis of a comparison of the output energy from the gates, noise in the energy measurement and comparator circuits can result in incorrect selection. In the next sub-section the bit-error probability arising from timing jitter is analysed in detail under the assumption that the gate selector is perfectly noise-free and has

infinitesimal resolution. In the succeeding sub-section the effect of errors in the operation of a realistic gate selector is considered.

3.2.2 Bit-error probability when using a perfect gate selector

As explained in section 3.1, the bit rate of the incoming data packet and the repetition frequency of the pulse train at the A input to the gate are assumed to be exactly equal, and the phase difference θ between them is taken to be constant over the duration of the packet. However, as explained above, the arrival time of an individual packet data bits is subject to fluctuations. Figure 7 shows the probability density function of the arrival time of the i th data bit at the gates, relative to the clock signal input to gate G1. We have taken the arrival time to be a normally distributed random variable with standard deviation σ (usually referred to as the rms jitter). The mean value of the distribution occurs $\theta T / 2\pi$ before a clock pulse. Purely for illustration, in Figure 7 the width of the probability density distribution has been exaggerated so that the various regions can be clearly seen (as we shall see later, the regenerator would normally be used in a regime in which the bit-arrival jitter σ is a much smaller fraction of the bit period T). If the actual arrival time of the i th data bit is within the unshaded region defined by $t_i - W < t < t_i$ (where t_i is the time of the i th clock pulse) and if the bit is a 'mark', the gate window will open and allow the i th clock pulse to be transmitted correctly thus regenerating the data 'mark'. If the arrival time of the i th data bit falls outside this unshaded region, then the i th clock pulse is not correctly modulated. The probability that the actual arrival time falls within the unshaded region is given by

$$p(\psi) = \Phi\left(\frac{\psi}{\sigma/T}\right) - \Phi\left(\frac{\psi - W/T}{\sigma/T}\right) \quad (5)$$

with $\psi = \psi_1$, where

$$25 \quad \psi_1 = \theta / 2\pi \quad (6)$$

$$\Phi(z) = \int_{-\infty}^z \phi(x) dx \quad (7)$$

$$\text{and } \phi(z) = \exp(-z^2/2) / \sqrt{2\pi} . \quad (8)$$

If the actual arrival time of the i th data bit is within the heavily shaded region defined by $t < t_{i-1} = t_i - T$ and if the bit is a 'mark', the gate window will open so as to allow the $(i-1)$ th clock pulse to be transmitted, with the possibility of

producing an error. The probability that the actual arrival time falls within the heavily shaded region is given by

$$q(\psi) = \Phi\left(\frac{\psi - 1}{\sigma/T}\right) \quad (9)$$

with $\psi = \psi_1$.

- 5 Similarly, if the actual arrival time of a the i th data bit is within the other heavily shaded region defined by $t > t_{i+1} - W = t_i + T - W$ and if the bit is a 'mark', the gate window will open so as to allow the $(i+1)$ th clock pulse to be transmitted, again with the possibility of producing an error. The probability that the actual arrival time falls within this region is given by

$$10 \quad r(\psi) = 1 - \Phi\left(\frac{\psi + 1 - W/T}{\sigma/T}\right) \quad (10)$$

with $\psi = \psi_1$.

At the gate G2, the clock pulses are arranged to arrive later in time by an amount $T/2$. In this case the expressions for the probabilities p , q and r are given by (5), (9) and (10), but with $\psi = \psi_2$ where

$$15 \quad \psi_2 = (\theta/2\pi + 1/2) \bmod 1 \quad (11)$$

For both gates, the bit error probability can be calculated by considering the various combinations of values of the j th data bit with its nearest neighbours, as presented in the table below (the small differences when $j=1$ or n , where there are n bits in the packet, are neglected). This table assumes that $0 < \psi_i < W/T$.

$(j-1)$ th data bit	j th data bit	$(j+1)$ th data bit	Bit-error probability in the j th bit position at the output from gate i
0	0	0	0
0	0	1	$q(\psi_i)/8$
0	1	0	$\{1 - p(\psi_i)\}/8$
0	1	1	$\{1 - p(\psi_i)\}\{1 - q(\psi_i)\}/8$
1	0	0	$r(\psi_i)/8$
1	0	1	$\{q(\psi_i) + r(\psi_i) - q(\psi_i)r(\psi_i)\}/8$
1	1	0	$\{1 - p(\psi_i)\}\{1 - r(\psi_i)\}/8$
1	1	1	$\{1 - p(\psi_i)\}\{1 - q(\psi_i)\}\{1 - r(\psi_i)\}/8$

Summing the right-hand column of the table, we obtain an expression for the total bit-error probability for the output from gate i , conditional on the phase angle θ :

$$\begin{aligned} B(\psi) &= \{1 - p(\psi)[\{1 - q(\psi)/2][1 - r(\psi)/2]\}\}/2 \\ &\approx \{1 - p(\psi)\}/2 \end{aligned} \quad (12)$$

where $\psi = \psi_1$ or ψ_2 as appropriate.

- 5 Figure 8 is a plot of the expression (12) against the phase angle θ , taking as an example the values $\sigma = 0.018T$ and $W = 0.75T$ for the jitter and gate window width, respectively. It can be seen that, with these parameter values, for any value of θ in the full range $0 \leq \theta < 2\pi$ the bit-error probability for the output from at least one of the gates is less than 10^{-12} , and therefore this bit-error probability
10 10 can be achieved continuously by selecting the output from one of the gates suitably in each time slot. In this work we take 10^{-12} as the target bit-error probability (or the equivalent probability of 10^{-9} that a packet of 1000 bits will contain an error).

- There will now be described an analysis of the method of gate selection based on
15 15 a comparison of the integrated energies of the packets emerging from each gate, as described in section 3.2.1. As noted earlier, the total optical energy measured at the output of a gate, integrated over the duration of the packet, is in effect a measure of the number of 'marks' appearing in the regenerated data packet. Similar to the calculation of the bit error probability given above, the expected
20 20 value of the optical energy measured at the output of the i th gate can be calculated by considering the various combinations of values of the j th data bit with its nearest neighbours, as presented in the table below.

$(i-1)$ th data bit	i th data bit	$(i+1)$ th data bit	Probability of a 'mark' in the position of the j th clock pulse at the output from gate i
0	0	0	0
0	0	1	$q(\psi_i)/8$
0	1	0	$p(\psi_i)/8$
0	1	1	$\{p(\psi_i) + q(\psi_i) - p(\psi_i)q(\psi_i)\}/8$
1	0	0	$r(\psi_i)/8$
1	0	1	$\{q(\psi_i) + r(\psi_i) - q(\psi_i)r(\psi_i)\}/8$
1	1	0	$\{p(\psi_i) + r(\psi_i) - p(\psi_i)r(\psi_i)\}/8$

1	1	1	$\{p(\psi_i) + q(\psi_i) + r(\psi_i) - p(\psi_i)q(\psi_i) - p(\psi_i)r(\psi_i) - q(\psi_i)r(\psi_i) + p(\psi_i)q(\psi_i)r(\psi_i)\} / 8$
---	---	---	--

The expected value of the optical energy measured at the output of the i th gate, $\langle E_i \rangle$, integrated over the duration of the packet, is found by summing the probabilities in the table above:

$$\begin{aligned} \langle E_i \rangle &= \frac{p(\psi_i) + q(\psi_i) + r(\psi_i)}{2} - \frac{p(\psi_i)q(\psi_i) + p(\psi_i)r(\psi_i) + q(\psi_i)r(\psi_i)}{4} + \frac{p(\psi_i)q(\psi_i)r(\psi_i)}{8} \\ 5 &\approx \frac{p(\psi_i) + q(\psi_i) + r(\psi_i)}{2} \end{aligned} \quad (13)$$

Here $\langle E_i \rangle$ is normalised to the value nw , where n is the number of bits in the original packet and w is the energy of a pulse representing a 'mark' (a single pulse representing a bit of value 1) in the output. Then an appropriate strategy for selecting the gate output in each packet time slot is:

- 10 if $\langle E_i \rangle = \max\{\langle E_1 \rangle, \langle E_2 \rangle\}$ then select gate i . (14)

It is assumed, for the remainder of this sub-section, that whatever mechanism is used for selecting the appropriate gate output in each time slot, its operation is ideal; e.g. the energy measurement and comparator circuits used to make a selection on the basis of (14) have no systematic errors, are perfectly noise-free
 15 and have infinitesimal resolution. Figure 9 shows a plot of the conditional bit-error probability $B(\theta)$ against the phase angle θ calculated according to (12), using an ideal mechanism to select the output gate i , and assuming for example that $W/T = 0.75$. It can be seen that in the case where the jitter is $\sigma = 0.018T$, the maximum value of B is 10^{-12} (as expected from the curves plotted in Figure 8). In
 20 the other case shown, where σ/T takes the larger value 0.03, the conditional bit-error probability $B(\theta)$ reaches a maximum value of 8×10^{-6} for certain values of the phase angle θ , but there are other ranges of θ in which $B < 10^{-12}$. Figure 10 shows a plot of the maximum bit-error probability B that a packet may suffer, for any phase angle θ within the range $0 \leq \theta < 2\pi$. The results are plotted versus the
 25 bit-arrival jitter σ , again assuming for example $W/T = 0.75$.

When a packet travels across a network it may pass through a sequence of regenerators. The phase angle θ differs from one regenerator to the next in a random fashion, because in this invention the local oscillators at the regenerators are free-running and their phases are uncorrelated. Therefore, at each regenerator

that the packet encounters, θ is a continuous uniform random variable on the interval $0 \leq \theta < 2\pi$, i.e. $\theta/2\pi \sim U(0,1)$. We assume here that the rms jitter σ in the arrival time of the data bits is the same at each regenerator. It follows that the bit-error probability after passing through N regenerators is given by

- 5 $N\langle e \rangle = N\langle \langle e|\theta \rangle \rangle$, where e denotes a random variable that takes the value 1 when a bit error occurs and the value 0 otherwise, and $\langle e|\theta \rangle$ denotes $B(\theta)$ the bit-error probability, conditional on the phase angle θ , for a single regenerator. By inspection, it can be seen that $\langle e|\theta \rangle$ is a periodic, symmetrical function of θ that exercises two full periods on the interval $0 \leq \theta < 2\pi$. Moreover $\langle e|\theta \rangle$ is continuous
10 and differentiable on the interval $W/2T < \theta/2\pi < (2W/T+1)/4$, which represents one half-period. Therefore the expected value $\langle \langle e|\theta \rangle \rangle$ with $\theta/2\pi \sim U[0,1]$ is identical to the expected value $\langle \langle e|x \rangle \rangle$ with $x \sim U[W/2T, (2W/T+1)/4]$, where $\langle e|x \rangle$ is given by (12). The expected value $\langle x \rangle = (4W/T+1)/8$ and $\text{var } x = 1/(12 \cdot 4^2) = 1/192$. Now let $Y = \ln\langle e|\langle x \rangle + \Delta \rangle$ and perform a Taylor
15 expansion:

$$Y = \ln\langle e|\langle x \rangle \rangle + \Delta \frac{d}{dx} \ln\langle e|\langle x \rangle \rangle + \frac{\Delta^2}{2!} \frac{d^2}{dx^2} \ln\langle e|\langle x \rangle \rangle + \frac{\Delta^3}{3!} \frac{d^3}{dx^3} \ln\langle e|\langle x \rangle \rangle + \dots$$

Taking expectations,

$$\langle Y \rangle = \ln\langle e \rangle \approx \ln\langle e|\langle x \rangle \rangle + \frac{1}{384} \frac{d^2}{dx^2} \ln\langle e|\langle x \rangle \rangle \quad (15)$$

- since $\langle \Delta \rangle = 0$, $\langle \Delta^2 \rangle = \text{var } x$ and $\langle \Delta^3 \rangle = 0$. To evaluate (15) we use the following
20 relations:

$$\frac{d^2}{dx^2} \ln\langle e|x \rangle = \frac{1}{\langle e|x \rangle} \frac{d^2}{dx^2} \langle e|x \rangle - \frac{1}{\langle e|x \rangle^2} \left(\frac{d}{dx} \langle e|x \rangle \right)^2,$$

$$\langle e|x \rangle \approx \frac{1}{2} \left[1 - \Phi\left(\frac{x}{\sigma/T}\right) + \Phi\left(\frac{x-W/T}{\sigma/T}\right) \right],$$

$$\frac{d}{dx} \langle e|x \rangle \approx \frac{1}{2\sigma/T} \left[\phi\left(\frac{x-W/T}{\sigma/T}\right) - \phi\left(\frac{x}{\sigma/T}\right) \right]$$

and

$$25 \quad \frac{d^2}{dx^2} \langle e|x \rangle \approx \frac{1}{2(\sigma/T)^3} \left[x \phi\left(\frac{x}{\sigma/T}\right) - (x-W/T) \phi\left(\frac{x-W/T}{\sigma/T}\right) \right].$$

- Figure 11 is a plot of the expected bit-error probability for a given packet, per regenerator passed, according to (15). The curves are plotted against the rms jitter σ , for various values of the gate window width W . For example, with $W = 0.75T$ and $\sigma = 0.032T$, the expected bit-error probability per regenerator
- 5 passed is 10^{-15} , or an expected bit-error probability of 10^{-12} after passing through 1000 regenerators. These results indicate that, for a given jitter σ , the bit-error probability for a packet after passing many regenerators can be expected to be substantially less than the maximum value for a single regenerator, as shown in
- 10 Figure 10. However, this result is unduly optimistic. It ignores the fact that, although the phase θ of a data packet relative to the local oscillator at a given regenerator is a continuous uniform random variable, nevertheless the phases of successive packets arriving at the regenerator *from the same source* may be highly correlated. This is because the physical fluctuations that cause θ to vary (such as temperature and mechanical changes acting on an optical fibre cable) may occur
- 15 on a much slower time scale than the arrival time of packets. Therefore the regenerated packet stream may contain bursts (and possibly prolonged bursts) of errors, even though the expected bit-error probability according to (15) is very low. Therefore it is appropriate to take the more conservative estimates of jitter tolerance indicated by equation (12) and plotted in Figure 9
- 20 In fact, an even more serious limitation is encountered if we drop the assumption that the gate selection is ideal (i.e. no systematic errors, noise-free, and having infinitesimal resolution), and this situation is analysed next.
- 3.2.3 Bit-error probability when using imperfect gate selection**
- If the mechanism used to perform the gate selection in each time slot is imperfect,
- 25 then at certain times the incorrect gate will be selected, leading to an increase in the bit-error probability. In this sub-section the gate selection technique will be analysed based on comparison of packet energies, as described above, in the case that the energy measurement and comparator circuits are subject to noise and other imperfections.
- 30 In Figure 5, the detectors D1 and D2 followed by electronic integrators each provide a voltage, V_1 and V_2 respectively, in each time slot which is proportional to the energy of the regenerated packet emerging from gate G1 and G2, respectively, in that time slot. The comparator C produces a binary output signal according to whether or not V_1 exceeds V_2 , and this signal is used to control the operation of

the selector switch S. Ideally, if the energy of the regenerated packet emerging from gate G1 is greater or equal to the energy of the regenerated packet emerging from gate G2, then the comparator output is 0 causing the switch S to select the output from gate G1; and ideally if the energy of the regenerated packet emerging from gate G2 is greater or equal to the energy of the regenerated packet emerging from gate G1, then the comparator output is 1 causing the switch S to select the other output. However because of systematic and random errors, the energy measurement and comparator output are non-ideal. Here we assume that systematic errors are negligible, and model the random errors by assuming that,

5 for any given phase angle θ , the voltage $V_1 - V_2$ is a normally-distributed random variable with variance σ_c^2 . Here the normalisation is such that a voltage V_1 or V_2 equal to 1 represents the energy of a regenerated packet consisting of n 'mark' pulses, where n is the number of bits in the original packet, and the expected number of 'marks' in the original packet is 0.5. Therefore the probability

10 15 $P(G1)$ that gate G1 is selected (i.e. that the comparator output signal is 0) is given by

$$P(G1) = \Phi\left(\frac{V_1 - V_2}{\sigma_c}\right), \quad (16)$$

and the probability that gate G2 is selected (i.e. that the comparator output signal is 1) is given by

$$20 \quad P(G2) = 1 - P(G1) = 1 - \Phi\left(\frac{V_1 - V_2}{\sigma_c}\right). \quad (17)$$

For a given phase angle θ , the bit-error probability is

$$B = P(G1)B(\psi_1) + P(G2)B(\psi_2), \quad (18)$$

where ψ_1 , ψ_2 and $B(\psi)$ are given by (6), (11) and (12), respectively.

Figure 12 shows a plot of the bit-error probability B against the phase angle θ , calculated according to (18), and assuming for example that $W/T = 0.75$, $\sigma/T = 0.018$ and $\sigma_c/T = 0.001$. By comparing with Figure 9 it can be seen that the bit-error probability is very significantly increased because of the non-ideal gate selection.

Figure 13 shows the severe degradation in bit-error probability that occurs when 30 the gate selection is non-ideal, even if the standard deviation σ_c/T is very small. The figure shows a plot of the maximum value of the bit-error probability B for any

phase angle θ in the range $0 \leq \theta < 2\pi$, calculated according to (18), and plotted versus σ_c/T . Again the calculations assume for example that $W/T = 0.75$ and curves for various values of σ/T are shown. Notice, for example, that with $\sigma/T = 0.018$, a value of σ_c/T as small as 10^{-7} is sufficient to degrade the bit-error probability by 4 orders of magnitude (maximum bit-error probability of 10^{-8} , as compared to the value 10^{-12} shown in Figure 10 in the case of ideal gate selection with $W/T = 0.75$ and $\sigma/T = 0.018$).

This limitation in the usefulness of the two-gate design for the bit-asynchronous regenerator is overcome by the quad-gate design described below. Alternatively a triple gate regenerator might be used, although this delivers a lesser improvement in performance.

3.4. Quad-gate bit-asynchronous optical packet regenerator

3.4.1 Principle of operation

Figure 14 shows a quad-gate bit-asynchronous packet regenerator. The data bits in the incoming packet are used to control the opening of four gates, G1, G2, G3 and G4. A data bit with value 1 ('mark') causes each of the four gates to open for a fixed time duration (the gate window), otherwise the gates remain closed. It is preferable, though not essential, that the widths of the time window for the four gates are equal. The output from the local clock (a continuous free-running source of optical RZ pulses at a repetition frequency nominally equal to the packet bit rate $1/T$) is applied to the inputs of the four gates, with relative delays in steps of $T/4$ as shown in Figure 14. As with the dual-gate regenerator, to ensure that the clock pulses will be correctly modulated by at least one of the four gates whatever the value of θ , the window width W must lie within a particular range of values, which for the quad-gate regenerator is $T/4 < W < T$. As previously, the incoming data pulses and local clock pulses are represented here by delta impulses. (As described in section 3.4.4, when finite pulse widths are taken into account the acceptable range of window widths is somewhat narrower than $T/4 < W < T$).

As in the case of the dual-gate regenerator, we should note, with reference to Figure 14, that there is an alternative and equally valid configuration in which the relative delays in steps of $T/4$ are removed from the input ports *A* of the gates, and placed instead at the control ports *C*. The operation of the regenerator is very similar in this case, and the predicted performance described later is the same.

Again, in the analysis below the configuration will be assumed to be that in which the packet data bits are connected directly to the control ports of the gates, and the differential delays are in the connections to the input ports.

As in Figure 5, the technique for gate selection shown in Figure 14 is based on a comparison of the total optical energy emerging from each gate, integrated over the duration of the packet. In this case the optical switch S must select one of the outputs from the four gates. As shown in Figure 14 the 1x4 switch could consist of an arrangement of three 1x2 switches, S₁, S₂ and S₃. The controller C measures the outputs from detectors D₁, D₂, D₃ and D₄ in each time slot, and sets the selector switches accordingly using the strategy described by (20) below. Other practical details are similar to those described in section 3.2.1. In particular, the optical delays (labelled L , $L + \Delta T/4$, $L + \Delta T/2$ and $L + 3\Delta T/4$ in Figure 14) between the outputs of the gates and the selection switch S are used to allow sufficient time for the circuitry and switch S to operate before the packets arrive at the switch. Typically the delay L will be slightly less than one time slot in duration. Optionally, the optical delays between the outputs of the various gates and the switch S may be made to differ in steps of $T/4$, as shown in Figure 14. The purpose of doing this is to compensate for the delay steps at the inputs to the gates, thus equalising the delay of all optical paths from the clock source to the output of the selector switch S. As noted previously, this has the benefit that all the regenerated packets are then in precise bit synchronism with each other and with the local clock, so that the local clock may therefore be used as a continuous and regular source of pulses for use in subsequent digital optical processing stages. A further very important benefit of equalising the delay in this way is to reduce the problem of 'packet slippage' in a large network, as described in section 3.4.3.

Also as noted in section 3.2.1, an alternative method of selecting the most appropriate gate output in each time slot may be to perform a bit-error measurement on the whole or part of the packet that emerges from each gate.

In the next sub-section the bit-error probability arising from jitter in the time of arrival of the incoming data bits is analysed.

3.4.2 Bit-error probability

The bit-error probability for the output from gate i is given by (12) with $\psi = \psi_i$, where

09545283-092200

$$\psi_i = \left(\theta / 2\pi + \frac{i-1}{4} \right) \bmod 1, \text{ with } i = 1, 2, 3, 4. \quad (19)$$

Figure 15 is a plot of the bit-error probability according to (12) and (19), against the phase angle θ , taking as an example the values $\sigma = 0.036T$ and $W = 0.75T$ for the jitter and gate window width, respectively. It can be seen that, with these 5 parameter values, for any value of θ in the full range $0 \leq \theta < 2\pi$ the bit-error probability for the output from at least one of the gates is less than 10^{-12} .
 The expected value of the optical energy measured at the output of the i th gate, $\langle E_i \rangle$, is given by (13). An appropriate strategy for selecting the gate output in each packet time slot, based on a comparison of the energies of the packets 10 emerging from each gate, is:

$$\text{if } \langle E_j \rangle = \min \{ \langle E_i \rangle \} \text{ then select gate } (1 + (j+1)\bmod 4) \quad (20)$$

where $\{ \langle E_i \rangle \}$ denotes the set $\{ \langle E_1 \rangle, \langle E_2 \rangle, \langle E_3 \rangle, \langle E_4 \rangle \}$.

Figure 16 shows a plot of the bit-error probability B against the phase angle θ , calculated according to (12), using an ideal mechanism to select the output gate / 15 (such as the strategy (20) executed by ideal circuits without noise or impairments) and assuming for example that $W/T = 0.75$. It can be seen that a maximum bit-error probability of 10^{-12} is predicted when $\sigma/T = 0.036$. By comparing this result with Figure 10, it can be seen that, compared to the dual-gate regenerator, the quad-gate regenerator can tolerate a higher level of rms jitter sigma in the 20 arrival time of the incoming data bits. The reason for this is that with the quad-gate regenerator, the clock pulses that are modulated by the selected gate are positioned further in time from the edges of the gate window, thus reducing the probability that jitter in the arriving data may cause a clock pulse to fall outside the window resulting in a bit error in the regenerated packet.
 25 The effect of non-ideal selection of the gate output is now analysed. In Figure 14, the detectors D1 to D4 followed by electronic integrators each provide a voltage, V_1 to V_4 respectively, in each time slot which is proportional to the energy of the regenerated packet emerging from gate G1 to G4, respectively, in that time slot. The comparator circuit C produces an output signal which is used to control the 30 operation of the selector switch S. The ideal output of the comparator is as defined by (20). However because of systematic and random errors, the energy measurement and comparator output are non-ideal. We assume here that the

DRAFT COPY DRAFT COPY

systematic errors are negligible, and model the random errors due to noise, etc, as follows.

We assume that, for a given phase angle θ , the voltage difference $V_i - V_j$ (for all i and $j \neq i$) is a normally-distributed random variable with mean V_{ij} and variance

- 5 σ_C^2 . As before, the voltages are normalised so that $V_i = 1$ represents the energy of a regenerated packet emerging from gate G_i consisting of n 'mark' pulses, where n is the number of bits in the original packet, and the packet is assumed to have a mark:space ratio of 1. Then, with probability

$$P(V_j = \min \{V_i\}) = \prod_{i \neq j} \Phi(V_{ij} / \sigma_C),$$

- 10 it will be determined on the basis of a measurement of the voltages V_i ($i = 1, 2, 3, 4$) that $\langle E_j \rangle = \min \langle E_i \rangle$. Then, if the result of that determination is used in the strategy (20) to select the gate output in each time slot, the resulting bit-error probability is

$$B = \sum_j B(\psi_j) \prod_{i \neq j} \Phi(V_{ij} / \sigma_C), \quad (21)$$

- 15 where $J = 1 + (j+1) \bmod 4$, and $B(\psi)$ is given by (12).

- Figure 17 shows a plot of the bit-error probability B against the phase angle θ , calculated according to (21), and assuming for example that $W/T = 0.75$, $\sigma/T = 0.036$ and $\sigma_C/T = 0.001$ and 0.052 . In this example, only when the parameter σ_C/T becomes as large as ~5% is there a significant increase in bit-
20 error probability. This demonstrates that the quad-gate regenerator can show good tolerance to non-ideal operation of the gate-selection mechanism. This is in contrast with the dual-gate regenerator discussed earlier, where it was found that the bit-error probability is severely degraded when the gate selection is non-ideal, even if the standard deviation σ_C/T is very small (Figure 13).

- 25 Figure 18 shows a plot of the maximum value of the bit-error probability B for any phase angle θ in the range $0 \leq \theta < 2\pi$, calculated according to (21), and plotted versus W/T . The results shown are for $\sigma/T = 0.036$ and $\sigma_C/T = 0.001$ and 0.05 . It is found that, depending on the value of σ_C/T , the optimum value of W/T lies in the range approximately 0.7–0.85.

Other calculations based on (21) show that with σ_c/T less than 4%, and $\sigma/T \leq 0.032$, the maximum bit-error probability is less than 10^{-15} , regardless of the phase angle θ . Thus for a packet which passes through a sequence of 1000 regenerators of this type, the bit-error probability will be less than 10^{-12} .

- 5 The quad-gate regenerator, with gate selection based on packet energy comparisons, works as well as it does because of the good design of the gate selection strategy, given by (20). This strategy ensures that the clock pulse is generally positioned well inside the gate window of the selected gate (for example, with $W/T = 0.75$, the minimum separation between the clock pulse and the edge of the gate window is about $T/4$), and this reduces the probability of bit errors in the regenerated packet due to timing jitter in the incoming data.

The quad-gate regenerator is therefore a preferred embodiment of the invention, and preferably W/T is in the range 0.7–0.85.

15 3.4.3 Packet time-position slippage

- A side-effect of this method of bit-asynchronous regeneration is that it introduces a small non-deterministic delay; in other words, the time delay between the arrival of an incoming data packet and the emergence of the regenerated packet consists of a constant delay plus a small non-deterministic component. In the case of the
- 20 quad-gate regenerator, provided all the optical paths between the clock source and the output of the selector switch S are equalised (as shown in Figure 14), to good approximation the non-deterministic delay introduced by the regenerator can be considered to be bounded on the interval $(0, T/4)$. In the case of a single regenerator this small random time displacement of the regenerated packet may be
- 25 considered negligible because it is only a fraction of the bit period T . However a packet that travels across a network through a sequence of many regenerators may accumulate a significant net time displacement. A system error may occur if the accumulated displacement exceeds some specified limit of L bit periods (such as the width of the time guard bands before and after each data packet).
- 30 This accumulated displacement in time experienced by a packet is similar to the problem of a random walk, or the Brownian motion of a particle, in one dimension. At the i th regenerator the regenerated packet undergoes a non-deterministic time shift τ_i relative to the average value of delay. To a good approximation, τ_i can be

DOCUMENT EDITION 022600

considered to follow a continuous uniform random distribution on the interval $(-T/8, T/8)$, with probability density function $U(\tau_i) = 4/T$. After a packet has passed through Q identical regenerators (with uncorrelated phases), its accumulated displacement from the average delay value is $\Sigma \tau_i$, which has an
 5 expected value of zero and variance $QT^2/192$. The probability that the magnitude of the accumulated displacement exceeds the specified limit is therefore given by

$$P(|\Sigma \tau_i| > LT) = 2P\left(Z > \frac{L}{\sqrt{Q/192}}\right), \quad (22)$$

where $Z \sim N(0,1)$ is the standard normal random variable, and it is assumed that Q is large (i.e. $Q > \sim 50$). For small values of Q we can use the property that
 10 $P(|\Sigma \tau_i| > LT)$ is exactly zero for $Q < 8L$ (since a directed (non-random) walk consisting of $8L$ steps, each of length not more than $T/8$, cannot reach a distance greater than $L T$).

Figure 19 shows values for the probability that a packet suffers 'slippage' (ie. a time displacement greater than the specified limit $L T$), according to (22). The
 15 packet slippage probability is plotted versus Q , the number of regenerators passed, for various values of $L T$. For networks in which $Q \leq 100$, 5-bit guard bands ($L = 5$) surrounding the data packet time slot would be sufficient to ensure a packet-slippage probability of less than 10^{-9} , whereas in larger networks with $Q \approx 1000$ guard bands of at least 14 bit periods are required. For a packet size of
 20 length $n = 1000$ bits, say, these guard bands represent a modest overhead on the network throughput.

However, the analysis leading to (22) and the results shown in Figure 19 are idealised, because it was assumed that all four optical paths between the clock source and the output of the selector switch S are precisely equal. These results
 25 give the minimum theoretical guard-band width in the ideal case. In practice the optical path equalisation inside each regenerator is subject to errors (errors in fabrication and assembly of the optical waveguides and components, as well as fluctuations in path lengths due to environmental effects), and so the effects of these must be considered.

30 Let us assume that a network contains of a large number of quad-gate regenerators, as shown in Figure 19, in each of which the optical paths between

00000000000000000000000000000000

the clock source and the output of the selector switch S differ in delay time from the average value by an amount δ , where δ is a normally-distributed random variable, $\delta \sim N(0, \sigma_D)$. The standard deviation σ_D is a measure of the errors in equalising the delays of the optical paths. A packet that passes through the i th regenerator will undergo a non-deterministic time displacement d_i , relative to the average value of delay, given by $d_i = \delta_i + \tau_i$, where τ_i is a continuous uniform random variable on the interval $(-T/8, T/8)$. After a packet has passed through Q regenerators, its accumulated displacement from the average delay value is $\sum d_i$. The expected value of $\sum d_i$ is zero, and since δ_i and τ_i are independent variables,

5 The expected value of $\sum d_i$ is zero, and since δ_i and τ_i are independent variables,

10 $\text{var } \sum d_i = Q(T^2/192 + \sigma_D^2)$. The probability that the magnitude of the accumulated displacement exceeds the specified limit of L bit periods is therefore given by

$$P(|\sum d_i| > LT) = 2P\left(Z > \frac{L}{\sqrt{Q(1/192 + (\sigma_D/T)^2)}}\right). \quad (23)$$

The probabilities calculated according to (23) produce a graph identical to Figure 19, except that the values of LT shown in the legend should be multiplied by a factor $\sqrt{1+192(\sigma_D/T)^2}$. For example, if the tolerance in equalising the optical paths in the regenerators is $\sigma_D = T$, then to prevent packet slippage after 100 regenerators ($Q = 100$) a guard band of at least $5 \times \sqrt{1+192} = 70$ bit periods is required, and for $Q = 1000$ a guard band of at least $17 \times \sqrt{1+192} = 195$ bit periods is required. These examples show that the errors in equalising the optical paths in

15 the regenerators should be minimised (preferably $\sigma_D < T$) to keep the guard bands as small as possible to avoid a costly reduction in network throughput.

20 At a peak data rate of 100 Gbit/s ($T = 100$ ps), a tolerance of $\pm T$ on path delay represents a length tolerance of ± 2 mm in glass, or approximately ± 1 mm on a semiconductor substrate. If the optical components in the regenerator are in the

25 form of an integrated monolithic semiconductor device, or discrete semiconductor devices mounted on a motherboard with planar waveguide interconnections, or discrete semiconductor devices connected by short optical fibre waveguides, these tolerances can be readily achieved. If, however, the optical gates are based on a nonlinear optical fibre device such as a fibre loop mirror, as suggested earlier, the

30 length of fibre used in each loop mirror may be as great as 1 km in order for the required optical intensity of the input data bits to be acceptably low. To fabricate

such devices to within a length tolerance of ± 2 mm represents a technical challenge. A practical solution would be to incorporate within each fibre loop a mechanical device to stretch a portion of the fibre so as to bring the overall length to the required value, or alternatively to incorporate an adjustable air delay line for 5 the same purpose. In addition it would be beneficial to house all the fibre loops together in an environmentally-controlled package, incorporating temperature control. Assuming a typical expansion coefficient of $10^{-6}/^{\circ}\text{C}$ for optical fibre, the temperature control would be required to be within $\pm 1 ^{\circ}\text{C}$ to maintain the length to within ± 2 mm, for loop lengths of ~ 1 km.

10 3.5 k -gate bit-asynchronous optical packet regenerator

It will be apparent to those skilled in the art, that the packet regenerator may be realised using different numbers of gates (for example 5 gates, or 8 gates) in addition to the 2-gate and 4-gate examples described above. Some general relations for a bit-asynchronous optical packet regenerator with k gates ($k > 3$) are 15 as follows. The bit-error probability for the output from gate i is given by (12) with $\psi = \psi_i$, and the expected value of the optical energy measured at the output of the i th gate is given by (13), where

$$\psi_i = \left(\frac{\theta}{2\pi} + \frac{i-1}{k} \right) \bmod 1 \quad (24)$$

and $i = 1, 2, \dots, k$. An appropriate strategy for selecting the gate output in each 20 packet time slot, based on a comparison of the energies of the packets emerging from each gate, is:

$$\begin{aligned} \text{if } \langle E_j \rangle = \min \{ \langle E_i \rangle \} \text{ then select gate } & \left(1 + \left(j - 1 + \frac{k}{2} \right) \bmod k \right) \text{ if } k \text{ is even,} \\ \text{or select gate } & \left(1 + \left(j - 1 + \frac{k \pm 1}{2} \right) \bmod k \right) \text{ if } k \text{ is odd} \end{aligned} \quad (25)$$

where $\{ \langle E_i \rangle \}$ denotes the set $\{ \langle E_1 \rangle, \langle E_2 \rangle, \dots, \langle E_k \rangle \}$.

Single Gate Regenerator

25 In an alternative embodiment, the local clock pulse source is again continuously free-running, but requires only one gate to modulate the output of the clock pulse source so as to regenerate the packet.

The principle of this alternative approach is shown in Figure 21. In each time slot, the phase detector measures the phase angle θ between the free-running local

pulse source and the incoming packet. This information is used to shift by an appropriate amount the phase of the control signal that is applied to the gate. The effect of the phase shifter is that when a packet data bit of value 1 causes the gate window to open, the window is located as near as possible centrally over the 5 clock pulse (as depicted in Figure 21). The phase detector and phase shifter operate once in each time slot.

In the ideal case the phase detector and phase shifter operate without error and the gate window is located exactly centrally over the clock pulse. In that case the main source of bit errors in the regenerated packet is jitter in the arrival time of the 10 packet data bits. The analysis of bit errors arising from timing jitter in this regenerator is then similar to the analysis of the effects of timing jitter in an OTDM demultiplexer given by Jinno (IEEE Journal of Quantum Electronics, vol 30, no.12, pp. 2842-2853, 1994). In particular, the bit-error probability as a function of the ratio of the gate widow width W to the bit period T , for various values of the rms 15 jitter σ in the arrival time of the packet data bits, is as depicted in Figure 5 of Jinno (1994). Some results of that analysis are that the minimum bit-error probability is obtained when W is equal to T , and also σ must be less than 0.071 T to ensure that the bit-error probability is less than 10^{-12} .

In a more realistic case, noise and other imperfections in the phase detector and 20 phase shifter cause systematic and random errors in the position of the gate window relative to the clock pulse. We neglect systematic errors and assume here that, as a result of the random errors, the position of the gate window is a normally distributed random variable with standard deviation σ_w . In that case, the bit-error analysis is again similar to that of Jinno (1994), and the results shown in 25 Figure 5 of Jinno can be used, except with the parameter σ replaced by $\sqrt{\sigma^2 + \sigma_w^2}$. Again it is found that the minimum bit-error probability is obtained when W is equal to T , and $\sqrt{\sigma^2 + \sigma_w^2}$ must be less than 0.071 T to ensure that the bit-error probability is less than 10^{-12} .

Figure 22 shows a possible embodiment of this alternative version of the bit- 30 asynchronous regenerator, and it is assumed that the bit rate is 100 Gbit/s. The incoming packet (at wavelength λ_m) is first passed through an optical input stage with slowly-responding automatic level control (such as an erbium-doped fibre amplifier), and is then split into two paths, one leading to the phase detector and

the other to the phase shifter. The local clock pulse source is a mode-locked ring fibre laser producing ~2 ps pulses at the wavelength λ_c . The output of this clock source is also split into two paths, one leading to the phase detector and the other to the optical gate. At the input to the phase detector, the clock pulses are 5 broadened to ~10 ps, for example by passing them through the optical bandpass filter F1 (as shown in Figure 22) or by dispersion in fibre or in a chirped grating. The phase detector could be based on four-wave mixing in a semiconductor optical amplifier SOA FWM (as described for example by O Kanatani, S Kawanashi and M Sarawutari in Electronics Letters, vol.30, no.10, p.807, 1994). The output from 10 the SOA FWM is isolated using the optical bandpass filter F2, and then detected. The electronic processing stage (which could include a low noise, high linearity sample-and-hold gate triggered by the global packet-level clock) measures the photodetector output voltage immediately after the arrival of the packet in each time slot, the measured voltage being given approximately by $A + B \cos\theta$, where A 15 and B are constants and θ is the phase difference between the clock and the incoming packet bits. In the example implementation of the phase shifter, this signal is used to control the wavelength λ_{cw} of the continuous-wave distributed-feedback laser DFB in the phase shifter section. The output of the DFB laser is connected to the input of a date-driven optical switch denoted UNI1. This device 20 could be the ultrafast nonlinear interferometer switch described by Hall and Rauschenbach (paper PD5, Proceedings of Conference on Optical Fiber Communication OFC'98, published by the Optical Society of America, February 1998), which has been shown to operate at a speed of 100 Gbit/s, although any data driven optical switch capable of producing approximately square switching 25 windows with negligible variation in throughput delay would be suitable. The control signal to UNI1 is the input data packet. The action of UNI1 is therefore to shift the wavelength of the incoming packet from λ_m to the controlled wavelength λ_{cw} . This packet with shifted wavelength is then isolated using optical filter F3, and passed through a dispersive optical delay line which imparts a time delay 30 which depends on wavelength. This dispersive delay line could be a fibre grating device or, as shown in Figure 22, a length of optical fibre, such as a length of optical fibre of the type manufactured primarily for use in dispersion compensation. The required minimum amount of phase shift of the control pulses is ±5 ps (i.e.

05072000 05072000 05072000

$\pm 0.5T$, where T is the bit period). Therefore, if for example the length of dispersion compensating fibre is 20m and the fibre dispersion is 100 ps/nm/km, the required shift in the wavelength λ_{cw} of the DFB laser is ± 2.5 nm, and this must be accomplished within the duration of the guard band between packets (say 5 ~ 5 ns). Alternatively, instead of controlling the wavelength of a single DFB laser, as shown in Figure 22, the control loop may contain means to select one of a number of lasers each having a different fixed wavelength. Because the optical fibre used as a dispersive delay line is a long device subject to changes in path length due to environmental factors (temperature, strain, etc), it is convenient as 10 in Figure 22 to pass the local clock pulses over the same fibre. This ensures that there is a negligibly small relative change in delay for the clock and control pulses. Alternatively, the error signal could be used to select one of a discrete number of 15 optical paths.

A further alternative phase shifter comprises means to select one from a 15 number of optical delay lines each having a different fixed delay. The delay lines could consist of a silicon wafer on which is fabricated a silica-on-silicon planar lightwave circuit. This circuit may be integrated in hybrid fashion with an array of discrete or integrated semiconductor optical switching devices, such as 20 semiconductor optical amplifiers or electroabsorption modulators. In this case the electronic processing stage selects the appropriate delay line by switching on or off the appropriate semiconductor optical switching devices.

The resultant optical data bits, suitably phase shifted, are then used as the control pulses in the optical gate. The gate, denoted UNI2 in Figure 22, which is controlled by the phase-shifted packet data bits, is used to modulate the locally 25 generated clock pulses so as to produce a regenerated packet, synchronous with the local clock. Again the device UNI2 could be the ultrafast nonlinear interferometer switch described by Hall and Rauschenbach. In the case that the regenerator is receiving inputs from a multiple number of sources, if the embodiment shown in Figure 22 is employed, the same length of dispersion- 30 compensating fibre should be used for all inputs to provide appropriate phase shifts, so that all the regenerated packets are bit-synchronous.

Figure 23 shows an example of how the bit-asynchronous regenerator may be used in an optical network. Three sources of packets are depicted (sources A, B and C), each of which have independent, uncorrelated clocks. Merely for clarity, in

Figure 23 the packets that have originated from source A, B or C are coloured white, black or shaded, respectively. The link from the output of each source carries only packets from that source, and therefore those packets are bit-synchronous with the clock in the source. This means that if a regenerator is required in the link, it may be a bit-synchronous type, similar for example to that depicted in Figure 2 . By suitable adjustment of the transmitted power at the source, the power levels in any optical amplifiers used in the link, and also the power levels at any synchronous regenerators used in the link, the bits in the packets arriving at the input of a routing node (such as routing node 1) may conveniently have an intensity at an appropriately-defined standard 'digital' level (e.g. of the correct intensity to perform complete switching in the optical gate or gates used in the bit-asynchronous packet regenerator AR in the switching node). The inputs to the switching nodes will, in general, be bit-asynchronous. Thus, for example in Figure 23, the packets that arrive at switching node 1 having originated from sources A and B (i.e. 'black' and 'white' packets), are bit-asynchronous because the clocks in sources A and B are uncorrelated. Each input to the routing node may pass through a bit-asynchronous packet regenerator AR, and each of these regenerators in a switching node share the same local optical clock pulse source. Each output from the routing switch may carry packets that have originated from more than one source. For example, in Figure 23, an output link from routing node 1 may contain packets that have originated from sources A and B (i.e. 'black' and 'white' packets), but the action of the bit-asynchronous packet regenerators AR in the switching node is such that all the packets carried on this output link are now in bit-synchronism with the local clock in routing node 1, despite their different sources. Therefore if a regenerator is required in this output link, it may be a bit-synchronous type. In Figure 23, the inputs to routing node 2 are from routing node 1 and from source C, and these inputs are bit-asynchronous. Each of these inputs pass through bit-asynchronous regenerators AR that share the same local optical clock pulse source, so that an output from the routing node 2 may contain packets that originate from various sources, including sources A, B and C, but all the packets on the output links are again in bit-synchronism with the local clock in routing node 2.

By using the bit-asynchronous packet regenerators in the switching nodes, we have shown that it is possible to design a network in which each individual link

D00160 - 28249456

carries packets which are bit-synchronous (i.e. which share the same bit-level clock) and which share a standard power level. This allows simpler designs of regenerators in the links, and also removes the need for packet-by-packet power level equalisation. However, by regenerating the packets at each node in bit-
5 asynchronous fashion, this entirely eliminates the need to maintain bit-level synchronism between the different links and routing nodes throughout the network - and thus we have eliminated the major architectural limitation of synchronous OTDM networks.

There is described above with reference to Figure 21, an alternative
10 approach to the bit-asynchronous packet regenerator, in which the local clock pulse source is again continuously free-running, but which requires only one gate to modulate the output of the clock pulse source so as to regenerate the packet. A measurement of the phase θ was used to control a phase shifter acting on the control signal to the optical gate. This alternative approach is, in effect, a feed-
15 forward, open loop control system. There is also described, in the immediately preceding section, the use of a regenerator in a network in which incoming packets at the regenerator are bit-synchronous (i.e. they all share the same bit-level clock). Therefore in this case the phase θ is only very slowly varying from one packet to the next, and this allows another arrangement for the control of the
20 phase shifter, as will be described with reference to Figures 24 and 25.

Figure 24 shows the principle of this alternative approach. In this case the phase detector measures the phase angle θ between the free-running local pulse source and the control signal applied to the optical gate (i.e. the packet data bits after the phase shifter). The measured phase angle is then used to control the
25 phase shifter. The control signal to the phase shifter may be either an analogue or digital signal (preferably a digital signal at the packet level), and this control signal would be gated at the rate of once per packet time slot. In contrast with the arrangement shown in Figure 21, this is a feed-back, closed loop control system. This has the advantage that the system is free of systematic errors and drift, even
30 if the phase detector is nonlinear. Furthermore, drifts in the phase shifter are automatically compensated for by the closed loop (because the phase shifter is inside the feedback path). Because, in a practical system, the feedback delay may be greater than the duration of a packet time slot, this closed-loop control system will not be sufficiently fast acting to track substantial variations in phase θ from

packet to packet. However this is not an important limitation in certain important network applications, as described above.

Figure 25 shows an example embodiment of this alternative form of the bit-asynchronous packet regenerator. The various designated components are as
5 described previously for Figure 22.

The description above includes a discussion of the allowable amount of frequency difference between the clock at the packet source and the clock in the asynchronous regenerator. The present embodiment of the asynchronous regenerator using a feed-back, closed-loop arrangement may impose a further
10 restriction on the amount of frequency difference that can be tolerated. It is necessary that the frequency difference between the bit rate of the incoming packet and the full-rate optical clock source in the regenerator is significantly smaller than the effective bandwidth of the control loop (including the electronic bandwidth, the feedback delay and the speed of response of the phase shifter).
15 Following normal engineering practice, the frequency offset should be at least an order of magnitude smaller than the effective bandwidth of the control loop. For example, if the determining factor for the control-loop bandwidth is an electronic bandwidth of 10 kHz, then the magnitude of the frequency offset $f_R - f_S$ should be no greater than 100 Hz (where f_S and f_R are the frequencies of the microwave
20 oscillators depicted in Figure 4, and it is assumed that $M_S = M_R = 10$).

In an alternative embodiment, the phase shifter shown in Figure 25, consisting of DFB laser, switching device UNI1, optical filter F3 and dispersion-compensating fibre, could be replaced by a variable optical delay line. Since, in the network scenario discussed above, and illustrated by Figure 23, the packets
25 arriving at the regenerator may be in bit-synchronism. Therefore the bit-asynchronous regenerator merely needs to track the relatively slow variations in the phase difference between the incoming packets and the local clock, rather than abrupt packet-to-packet phase variations. The control loop may therefore be relatively slow acting (much slower than on a packet-by-packet basis). However,
30 the control loop bandwidth should not be so low as to restrict the amount of frequency offset that can be accommodated. Suppose, for example, that the delay line is a variable motor-controlled device, such as a motor-driven fibre stretcher, which is capable of changing the value of the optical delay at a maximum rate of

03524582 0222000

100 ps per second. This corresponds to 10 bit periods per second at 100 Gbit/s, for example, and so the maximum allowable frequency offset would be an order of magnitude less than that (to ensure the effective control loop bandwidth is at least 10 times faster than the fastest variations in the signal to be controlled), i.e. ~1 5 Hz, which is a severe restriction. Therefore a motor-controlled phase shifter may not have sufficient speed of response for this application. Another type of variable optical delay line is a fibre stretcher consisting of a length of fibre coiled tightly around a piezo-electric drum. This type of stretcher is capable of 100 μm length change at 20 kHz, or approximately 1 ps delay change in 50 ms. In order to 10 achieve a range of ± 5 ps (i.e. ± 0.5 bit periods at 100 Gbit/s), a cascade of piezo-electric drum stretcher units (10-20 units) would allow a frequency offset of a few kHz between the local and distant clocks.